

CONNECTION VERIFICATION FOR OPTICAL SWITCHESFIELD OF THE INVENTION

5

The present invention relates generally to optical switches and, more particularly, to systems for verifying connections through such switches.

10 BACKGROUND OF THE INVENTION

As the density of emerging high-capacity WDM systems increases, so too does the probability with which errors can be made when switching individual optical signals. Examples
15 of what may cause an erroneous or lost connection include stuck or failed switch elements in the switching core of an optical switch, hardware or software failures causing incorrect switch path instructions to be received by the switching core from a switch controller, and human error
20 (e.g., a mis-connected fiber interconnect into or between bays of switching equipment).

Given the high line rates currently used in WDM networks and the even higher line rates contemplated for use in the
25 foreseeable future, it is clear that erroneous or lost connections can and will have a very severe negative impact on quality of service by causing the loss of large amounts of information. It is therefore of prime importance to check not only whether connections established by the switching
30 core correspond to the connections specified by the connection map stored in the switch controller, but also to check which mis-connections may have taken place.

At the same time, it is important to make the connection verification process as generic and non-intrusive as possible so that constraints are not placed on the traffic bit rates and protocols, so that the effect that the connection verification procedure has on the quality and strength of the optical signals leaving the switch is limited and so that traffic security is not compromised. Moreover, it would be of interest if the connection verification system were to permit the use of simple, low-cost electronics and electro-optics for the verification function in order to limit the cost and component count of the additional hardware and software required to verify the connections. However, these requirements have yet to be met by existing connection verification techniques.

SUMMARY OF THE INVENTION

The present invention provides substantially non-intrusive connection verification for an optical switch, achieved by correlating the low-frequency contents of the input and switched signals. The results of the correlation process are compared against a connection map to determine whether the switch has operated correctly and to identify, if possible, which mis-connections have taken place.

Accordingly, the invention may be summarized according to a first broad aspect as a system for verifying connections established through a switching unit adapted to receive a plurality of input signals and output a plurality of switched signals. The system includes a selection unit for controllably admitting individual ones of the input signals and individual ones of the switched signals and a verification unit connected to the selection unit, for controlling operation of the selection unit as a function of

a connection map and performing relative-delay-dependent signal processing operations on the signals admitted by the selection unit so as to identify connections established through the switching unit and determine their consistency
5 with the connection map.

The invention may be summarized according to a third broad aspect as a method of validating connections established through a switching unit adapted to receive a plurality of
10 input signals and output a plurality of switched signals. The method includes selecting one of the input signals; on the basis of a connection map, identifying a particular one of the switched signals as expected to be correlated with the selected input signal; determining a level of correlation or
15 anti-correlation between the selected input signal and the switched signal expected to be correlated with the selected input signal; and if the level of correlation is significant or the level of anti-correlation is insignificant, concluding that the connection involving the selected input signal is
20 consistent with the connection map.

The invention may also be broadly summarized as computer-readable media tangibly embodying a program of instructions executable by a computer to perform the above method. The
25 invention may also be broadly summarized as at least one computer programmed to execute the above method.

According to another broad aspect, the invention may be summarized as a switch for optical signals, including a
30 switching core for switching a plurality of input optical signals as a function of a connection map and outputting a plurality of switched optical signals and a connection verification system connected to the switching core, for correlating the input optical signals with the switched

optical signals so as to determine the consistency of the connections established through the switching core with the connection map.

5 The invention may be summarized according to yet another broad aspect as a system for correlating a first sample stream with a second sample stream, including a first delay line for receiving the first sample stream, comprising a plurality of taps interspersed by delay elements, a second
10 delay line for receiving the second sample stream, comprising a plurality of taps interspersed by delay elements, a plurality of anti-correlators, each having a first input connected to a tap in the first delay line and a second input connected to a tap in the second delay line, each correlator
15 being adapted to produce an anti-correlation value at a distinct relative delay, and a minimum detector connected to the anti-correlators, for selecting the least among the anti-correlation values produced by the anti-correlators and providing the result to a controller.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other aspects and features of the present invention will now become apparent to those of ordinary skill in the
25 art upon review of the following description of specific embodiments of the invention in conjunction with the accompanying drawings, in which:

Fig. 1 is a block diagram of a WDM photonic switch;

30 Fig. 2 illustrates mappings occurring within the photonic switch of Fig. 1;

Fig. 3 is a block diagram of a subsystem used for the verification of a first subset of connections through the photonic switch of Fig. 1;

5 Fig. 4 is a block diagram of a front end for use in the subsystem of Fig. 3;

Fig. 5 is a block diagram of a connection verification system for use in the subsystem of Fig. 3;

10 Figs. 6A-6C are flowcharts illustrating operational steps in a connection verification algorithm executed by a controller in the connection verification system of Fig. 5; and

15 Fig. 7 is a block functional diagram of a differential correlator in the connection verification system of Fig. 5.

Figs. 8-11 are block diagrams of various alternative implementations of the front end of Fig. 4;

20 Fig. 12 is a block diagram of a subsystem used for the verification of a second subset of connections through the photonic switch of Fig. 1;

25 Fig. 12A is a block diagram of an alternative embodiment of the subsystem in Fig. 12;

Fig. 13 is a block diagram of an input signal front end for use in the subsystem of Fig. 12;

30 Fig. 14 is a block diagram of a subsystem used for the verification of a third subset of connections through the photonic switch of Fig. 1;

Fig. 14A is a block diagram of an alternative embodiment of the subsystem in Fig. 14;

Fig. 15 is a block diagram of a subsystem used for the verification of a fourth subset of connections through the photonic switch of Fig. 1;

Fig. 15A is a block diagram of an alternative embodiment of the subsystem in Fig. 15;

10

Fig. 16 is a block diagram of a subsystem used for the verification of a fifth subset of connections through the photonic switch of Fig. 1;

Fig. 16A is a block diagram of an alternative embodiment of the subsystem in Fig. 16;

Fig. 17 is a block diagram of a subsystem used for the verification of a sixth subset of connections through the photonic switch of Fig. 1;

Fig. 17A is a block diagram of an alternative embodiment of the subsystem in Fig. 17;

Fig. 18 is a block diagram of an alternative implementation of the connection verification system of Fig. 5.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

As will be described in further detail herein below, the connection verification system of the present invention correlates the low-frequency content of a set of input single-carrier optical signals with the low-frequency content of a set of switched single-carrier optical signals. For

best performance of the connection verification system, the low-frequency content (representing the first, say, 10-100 MHz of signal spectrum, depending on the design) of the various input signals arriving at the switch should be sufficiently unique and/or randomized so that the cross-correlation of the low-frequency content of any pair of distinct input signals is relatively weak compared to each signal's auto-correlation. This condition is met for most forms of scrambled binary information transmission including those commonly in use with existing fiber systems, such as SONET, SDH, PDH, Gigabit Ethernet, 10GE, as well as direct IP-over-wavelength, and wavelength-wrapper-encapsulated signals.

Other examples of signals suitable for use with the present invention include signals having a low-frequency component that is encoded using Walsh codes or other spread spectrum techniques. Suitable signals further include those designed in accordance with United States Patent Application Serial No. 09/648,767 to Graves et al., filed on August 28, 2000, entitled "Method, System and Signal for Carrying Overhead Information in a Transport Network Employing Photonic Switching Nodes", assigned to the assignee of the present invention and hereby incorporated by reference herein. The signals described in United States Patent Application Serial No. 09/648,767 have a controllable low-frequency content (which can be made unique for each fiber and wavelength combination), while retaining enough high-frequency content to allow synchronization of downstream network elements to be maintained.

Also, while the connection verification system of the present invention is applicable to virtually any type of switch, it is particularly suitable for use with a "photonic" switch. A

"photonic" switch is a switch for optical signals (i.e., an optical switch) where the bulk of the traffic paths through the switch node are entirely implemented in the optical domain, i.e., without the need for converting optical signals into the electrical domain. An example of a photonic switch is shown in Fig. 1 and is described in greater detail in United States Patent Application Serial No. 09/511,065 to Graves et al., filed on February 23, 2000, entitled "Switch for Optical Signals", assigned to the assignee of the present invention and hereby incorporated by reference herein.

The functionality of a photonic switch is to switch or interchange individual single-wavelength modulated optical carriers within a series of input multi-carrier WDM feeds and to re-multiplex the resultant new combinations of single-carrier optical signals into new multi-carrier WDM feeds, having first wavelength-shifted any single-carrier optical signals that needed this function in order to permit onward propagation through the network. As such, a photonic switch is optimized for use in WDM line systems.

Specifically, and with reference to Fig. 1, a photonic switch 100 comprises a switching core which, in one embodiment, includes a plurality of optical switch matrices 110A...110M (one for each of M optical wavelengths in the system), as well as a wavelength converting switch 120. The switching core provides controllable switching of single-carrier optical signals received from a plurality of wavelength division demultiplexing (WDD) devices 130A...130N, which demultiplex the incoming multi-carrier WDM feeds into individual single-carrier optical signals. Switched single-carrier optical signals emerge from the switching core and are provided to a plurality of wavelength division multiplexing (WDM) devices 140A...140N, which regroup the

switched single-carrier optical signals into multi-carrier WDM feeds for onward propagation to the next network node.

Each of the optical switch matrices 110A...110M has a total of
5 N+K input ports and N+K output ports. For a given one of the
optical switch matrices 110A...110M, each of N input ports is
connected to the like-wavelength output port of a respective
one of the WDD devices 130A...130N, while the remaining K input
ports are connected to output ports of the wavelength
10 converting switch 120. In an analogous fashion, each of N of
output ports of each optical switch matrix is connected to
the like-wavelength input port of a respective one of the WDM
devices 140A...140N, while the remaining K output ports are
connected to input ports of the wavelength converting switch
15 120.

Each of the optical switch matrices 110A...110M can be a Micro-
Electro-Mechanical System (MEMS) as described in an article
entitled "Free-Space Micromachined Optical-Switching
20 Technologies and Architectures" by Lih Y. Lin of AT&T Labs-
Research, presented during OFC99 Session W14-1 on February
24, 1999 and hereby incorporated by reference herein. As is
described in the aforementioned article, a MEMS comprises a
set of mirrors that are arranged in geometrical relationship
25 with its input and output ports in such a way that incoming
light from any input port can be diverted to any output port
by moving an appropriate one of the mirrors.

For the photonic switch 100 of Fig. 1, the moving (e.g.,
30 raising and lowering) of mirrors in the optical switch
matrices 110A...110M could be performed under control of a
switch controller 150. Specifically, the switch controller
150 can supply each of the optical switch matrices 110A...110M
with a connection map defining the desired switching

behaviour of that optical switch matrix. With additional reference to Fig. 2, the connection map supplied to the λ^{th} optical switch matrix 110 λ can take the form of an $(N+K) \times (N+K)$ matrix, which may be represented as follows:

5

output→		1	2	...	N-1	N	N+1	...	N+K
↓input									
	1	-----							
	2								
10	...			$[A_\lambda]$				$[C_\lambda]$	
	N-1								
	N								
		-----					+-----		
	N+1								
15	...			$[B_\lambda]$				$[NULL]$	
	N+K	-----							

where

20 $[A_\lambda]$ is a sparse matrix defining the desired mapping between input ports 1...N and output ports 1...N, where a "1" in position (x,y) signifies that the optical signal arriving from WDD device 130x and occupying the λ^{th} wavelength is to be switched directly to WDM device 140y without wavelength conversion;

25 $[B_\lambda]$ is a sparse matrix defining the desired mapping between input ports N+1...N+K and output ports 1...N, where a "1" in position (x,y) signifies that the optical signal arriving from the $(x-N)^{\text{th}}$ output of the wavelength converting switch 120 (for the λ^{th} wavelength) is to be switched to WDM device 140y; and

30 $[C_\lambda]$ is a sparse matrix defining the desired mapping between input ports 1...N and output ports N+1...N+K,

where a "1" in position (x,y) signifies that the optical signal arriving from WDD device 130x is to be switched to the (y-N)th input of the wavelength converting switch 120 (for the λ^{th} wavelength).

5

Thus, there should be no more than a single "1" in any row or column of $[A_\lambda]$, $[B_\lambda]$ or $[C_\lambda]$. Furthermore, there should be no "1" in any row of $[A_\lambda]$ where there is a "1" in that row of $[C_\lambda]$ and there should also be no "1" in any column of $[A_\lambda]$ where there is a "1" in that column of $[B_\lambda]$.

For its part, the wavelength converting switch 120 receives K single-carrier optical signals from each of the M optical switch matrices 110A...110M and outputs K single-carrier optical signals back to each of the M optical switch matrices 110A...110M. The wavelength converting switch 120 may also accept a total of R "add carriers" on a respective plurality of optical add paths 180A...180R and similarly may output a total of R "drop carriers" on a respective plurality of optical drop paths 170A...170R. Of course, it is within the scope to have different numbers of add carriers and drop carriers.

The wavelength converting switch 120 is equipped with circuitry for converting the received single-carrier optical signals into electrical form, electrically switching the electrical signals and then modulating each switched electrical signal with an optical source. The optical source may be fixed or tuned to the wavelength associated with a specific optical switch matrix in the switching core. Wavelength conversion is particularly useful when an input wavelength is already in use along the fiber path leading to a specific destination WDM device or group of WDM devices

140A...140N that lead to the required destination via ongoing output optical fibers.

The switching activity of the wavelength converting switch 120 is controlled by the switch controller 150 as a function of a connection map. With reference to Fig. 2, the connection map supplied to the wavelength converting switch 120 can take the form of an $(M \cdot K + R) \times (M \cdot K + R)$ matrix, which may be represented as follows:

10

output→		1	2	...	M•K	M•K	M•K	...	M•K
↓input					-1		+1		+R
1		<hr/>							
2									
...				[D]				[F]	
M•K - 1									
M•K									
		<hr/>							
M•K + 1									
...				[E]				[NULL]	
M•K + R		<hr/>							

where

25 [D] is a sparse $M \cdot K \times M \cdot K$ matrix defining the desired mapping between outputs N+1 to N+K of all M optical switch matrices 110A...110M and inputs N+1 to N+K of all M optical switch matrices 110A...110M;

30 [E] is a sparse $R \times M \cdot K$ matrix defining the desired mapping between the R add carriers and inputs N+1 to N+K of all M optical switch matrices 110A...110M; and

[F] is a sparse $M \cdot K \times R$ matrix defining the desired mapping between outputs $N+1$ to $N+K$ of all M optical switch matrices 110A...110M and the R drop carriers.

5 From the above, it should be appreciated that the desired switching behaviour of the switching core will be an intricate function of mappings $[A_\lambda]$, $[B_\lambda]$, $[C_\lambda]$, $[D]$, $[E]$ and $[F]$ (where $\lambda \in \{A, \dots, M\}$ for mappings $[A_\lambda]$, $[B_\lambda]$ and $[C_\lambda]$). However, in the event of a malfunction or mis-connection, the
 10 switching core might not exhibit the desired switching behaviour. With the aim of obtaining more information as to the source of a malfunction or mis-connection, connection verification operations can be performed on various mapping subsets.

15 With reference to Fig. 3, there is shown a minimally intrusive subsystem for verifying whether each mapping in the following subset of mappings is being executed properly by the optical switch matrices 110A...110M:

$$20 \quad \{[A_\lambda], \lambda \in \{A, \dots, M\}\} \cup \{[C_{\lambda x}] \times [B_{\lambda y}], (\lambda x, \lambda y) \in \{A, \dots, M\}^2\}.$$

25 The first part of the above subset of mappings covers signals entering and exiting the photonic switch at the same wavelength using the same optical switch matrix.

The second part of the above subset of mappings includes
 30 composite mappings which cover the case where signals undergo wavelength conversion. It will be noted, however, that if there is an error in the way a particular mapping (either $[B_\star]$ or $[C_\star]$) is executed, connection verification performed using the subsystem of Fig. 3 will not allow precise

identification of the culprit as either [B*] or [C*] since it covers only composite mappings. To determine precisely which mapping is being erroneously executed requires that connection verification be performed using a different subsystem which is tailored towards verifying all mappings [B_{λy}] and [C_{λx}] individually and such subsystems are described later on with reference to Figs. 14 and 15, respectively.

Continuing with the description of the subsystem of Fig. 3, there is provided a first set of N optical splitters 310A...310N, one in the optical path of each WDM signal entering a respective one of the WDD devices 130A...130N in the photonic switch 100 of Fig. 1. Each of the N optical splitters 310A...310N diverts a small fraction of the corresponding WDM signal towards a common input signal front end 320. The input signal front end 320 provides a band-limited electrical signal to an input signal side of a connection verification system 330.

With additional reference to Fig. 4, there is shown a detailed block diagram of one of many embodiments of a front end 410 suitable for use with the present invention. The front end 410 could be the input signal front end 320 of Fig. 3. The front end 410 includes a set of N WDD devices 420A...420N, each of which is connected to a respective one of the optical splitters 310A...310N and has a respective set of M outputs, one for each of the M wavelengths in the system. Each output of each of the WDD devices 420A...420N is supplied to a respective one of a plurality of receivers 430.

Each of the receivers 430 is operable to convert the corresponding incoming single-carrier optical signal into a band-limited electrical signal. This reduced-bandwidth

signal will be used for correlation purposes in the connection verification system 330. Because of its reduced bandwidth, the signal received by the connection verification system 330 does not permit the recovery of a full-rate digital data stream and hence does not pose a security risk. Furthermore, since the optical receivers 430 only need exhibit enough bandwidth to pass the frequency spectral components required for the operation of 330, they can be of a relatively low, fixed bandwidth, independent of the traffic signal bandwidth.

Alternatively, each receiver 430 can have a wide electrical bandwidth, with the precise bandwidth of the signal to be fed to the connection verification system 330 being determined by one or more low-pass electrical filters (not shown) between the receivers 430 and the connection verification system 330. The electrical bandwidth of the receivers 430 (or of the receivers 430 plus the low-pass filters) thus sets a coarse upper bound on the bandwidth of the resulting signal that is used for correlation purposes within the connection verification system 330. A suitable electrical bandwidth for the receivers 430 will also allow for relatively inexpensive implementation of the front end 410. An example of a suitable electrical bandwidth for the receivers 430 is 100 MHz, although other higher and lower bandwidths can be used, depending on the operational requirements of the invention.

The front end 410 also comprises a set of N first M-way selectors 440A...440N, each of which receives the output of M respective receivers 430. Specifically, the set of first selectors 440A...440N is arranged so that first selector 440n receives those M electrical signals that correspond to the M optical signals provided by WDD device 420n for $n \in \{A, \dots, N\}$. Operation of the first set of selectors 440A...440N is

jointly controlled via a select line 542, which effectively selects the wavelength (denoted λ_x). The output of each of the N first selectors 440A...440N is provided to a respective input of a single second selector 450. Operation of the second selector 450 is controlled via another select line 544, which effectively selects the port (denoted p_x) of the input optical switch matrix at the selected wavelength (λ_x).

Thus, the output of the second selector 450 is a band-limited electrical signal provided to the connection verification system 330. This band-limited electrical signal is the one appearing at input port p_x of optical switch matrix 110 λ_x . It should be appreciated that control of the first selectors 440A...440N results in selection of the wavelength λ_x , while control of the second selector 450 results in selection of p_x , corresponding to the port or, equivalently, to the WDD device 420A...420N. Both free scanning and directed control of sequencing λ_x and p_x can be exerted.

Continuing with the description of the subsystem in Fig. 3, there is provided a second set of N optical splitters 340A...340N, one in the optical path of each switched WDM signal exiting a respective one of the WDM devices 140A...140N. Each of the N optical splitters 340A...340N diverts a small fraction of the corresponding switched WDM signal towards a common switched signal front end 350. The switched signal front end 350 could be identical to the front end 410 of Fig. 4, except that select lines 542 (specifying λ_x) and 544 (specifying port p_x or WDD device 130 p_x) would be replaced by select lines 543 (specifying wavelength λ_y) and 544 (specifying port p_y or WDM device 140 p_y).

Thus, the switched signal front end 350 provides a band-limited electrical signal to a switched signal side of the

connection verification system 330. This band-limited electrical signal is the one appearing at output port py of optical switch matrix 110λy.

5 Thus, it is seen that the connection verification system 330 receives a band-limited input signal from the input signal front end 320 (controlled by λx and px) and also receives a band-limited switched signal from the switched signal front end 350 (controlled by λy and py). The connection
10 verification system 330 receives the set of mappings [Aλ], [Bλ] and [Cλ] (λ ∈ {A, ..., M}) from the switch controller 150.

Through precise control of px, py, λx and λy, a controller within the connection verification system attempts to
15 determine whether the subset of mappings defined by:

$$\{[A_\lambda], \lambda \in \{A, \dots, M\}\} \cup$$

$$\{[C_{\lambda x}] \times [B_{\lambda y}], (\lambda x, \lambda y) \in \{A, \dots, M\}^2\}$$

20 has been accurately executed by the optical switch matrices 110A...110M. This will be described in greater detail in the course of a description of an example connection verification system.

25 Accordingly, reference is now made to Fig. 5, wherein is shown a functional block diagram of a connection verification system 330 in accordance with an embodiment of the present invention. The connection verification system 330 comprises
30 an input signal sampling module 530, which accepts the band-limited electrical signal consisting of the low frequency components of the time-varying frequency content of the digital bit stream under test from the input signal front end 320. The output of the sampling module 530 is an input

signal sample stream, which is provided to a first in a series of delay elements 550 forming an input signal tapped delay line. The delay elements 550 may be implemented as shift registers, for example.

5

In one embodiment of the connection verification system 330, the sampling module 530 may include a precision anti-alias filter and an analog-to-digital converter (ADC). Since the anti-alias filter will establish the frequency spectrum that will be digitized and used in the correlator, the sampling rate of the ADC and the bandwidth of the anti-alias filter should be selected such that the Nyquist criterion is satisfied. An anti-alias filter is not required if the electrical bandwidth of the receivers 430 in the input signal front end 320 is less than half of the sampling rate of the ADC in the sampling module 530. If such a filter is omitted, then the requirement for maintaining the same filtered bandwidth for the input and switched signals moves to the receivers 430.

20

The connection verification system 330 of Fig. 5 also comprises a switched signal sampling module 532. The switched signal sampling module 532 may be identical to the input signal sampling module 530 and as such may include a precision anti-alias filter and an ADC. The output of the switched signal sampling module 532 is a switched signal sample stream that is provided to a first in a series of delay elements 552 forming a switched signal tapped delay line.

30

The purpose of the two tapped delay lines containing delay elements 550, 552 is to allow a series of parallel correlations to be performed with different time "skew" between the input samples and the output samples. This is to

allow the different potential delays through the switch node to be taken into account during the correlation process. For example, for a propagation distance of 20 meters through the switch at about 200,000 km/sec (the approximate speed of light in glass), the delay which would need to be accommodated is approximately equal 100 nanoseconds.

Continuing with the description of Fig. 5, different pairs of samples, one from each of the two sample streams, are fed to respective ones of a plurality of differential correlators 560A...560E. This causes a distinct relative delay to exist between each different pair of samples fed to each of the differential correlators 560A...560E.

Specifically, differential correlator 560A is fed an undelayed sample from the input signal tapped delay line and a four times delayed sample from the switched signal tapped delay line (for a relative delay of -4), differential correlator 560B is fed a once delayed sample from the input signal tapped delay line and a thrice delayed sample from the switched signal tapped delay line (for a relative delay of -2), differential correlator 560C is fed a twice delayed sample from the input signal tapped delay line and a twice delayed sample from the switched signal tapped delay line (for a relative delay of 0), differential correlator 560D is fed a thrice delayed sample from the input signal tapped delay line and a once delayed sample from the switched signal tapped delay line (for a relative delay of +2) and differential correlator 560E is fed a four times delayed sample from the input signal tapped delay line and an undelayed sample from the switched signal tapped delay line (for a relative delay of +4).

It will be understood by those of ordinary skill in the art that in a unidirectional switch (one with "inputs" and "outputs"), only one set of delay elements may be required (e.g., the set of delay elements 550), since the input will typically always precede the output. However, both sets of delay elements will be required in the case of long optical paths from the optical splitters 310A...310N to the connection verification system 330 and also in the case where the photonic switch 100 is worked in a bidirectional fashion (in which case the inputs and outputs become interchangeable).

It should also be appreciated that although only four delay elements are shown in each of the tapped delay lines of Fig. 5, the actual number of delay elements 550 and 552 and the individual delay provided by each such delay element should be chosen as a function of the operational requirements of the invention. Thus, a wide variation in the individual delay of each delay element and in the total delay of each tapped delay line is within the scope of the present invention.

Nonetheless, the chances of obtaining a meaningful result are improved when the total delay of each tapped delay line is at least as long as the maximum lag through the switch (which is usually on the order of less than 100 nanoseconds but may be more).

Furthermore, since the cross-correlation of two signals is sensitive to the relative delay between them, it is preferable to ensure that the change in the relative delay, from one differential correlator to the next, is kept to within a predetermined upper bound.

For instance, it has been found that meaningful results can be obtained when the change in the relative delay between pairs of signals fed to adjacent differential correlators is less than about 15-30 degrees at the signals' highest frequency component. This is because the accuracy of the correlation process degrades significantly beyond a phase error of about 15-30 degrees between the signal components being correlated. For signals with a bandwidth of 15 MHz each, this corresponds to a requirement that the relative delay vary by less than 10 ns from one differential correlator to the next. In the connection verification system of Fig. 5, where the delay "resolution" is twice the delay produced by a delay element (in this case either -4, -2, 0, +2 or +4), each delay element would be required to provide a delay of at most 5 ns.

One possible implementation of a differential correlator is now described with additional reference to Fig. 7, wherein it is seen that the differential correlator accepts samples of an input signal from the input signal tapped delay line and samples of a switched signal from the switched signal tapped delay line. The differential correlator is also seen to accept a clock signal from a common external clock signal generator 590, as well as a counter signal from a common external clock cycle counter 580. The clock cycle counter 580 provides a running total of the clock cycles generated by the clock signal generator 590.

Samples of the input signal and the switched signal enter a subtractor 710, which may be of standard design and which is clocked by the clock signal received from the clock signal generator 590. The subtractor 710 computes the difference between the input and switched signal samples and provides the result to both inputs of a multiplier 720. Note that,

for two perfectly correlated signals, the output of the subtractor 710 would be zero, but for imperfectly or non-correlated samples the output would be non-zero, i.e., of either positive or negative polarity.

5

The multiplier 720, also clocked by the clock signal from the clock signal generator 590, computes the product of the samples at its two inputs. In this case, the product corresponds to the square of the sample at the output of the subtractor 710, which is a measure of how far apart are the two inputs to the subtractor 710. The output from 720 will always be positive for a case of a difference between the two inputs of the subtractor 710 irrespective of the polarity of that difference, but will be zero when those inputs are the same. The product generated at the output of the multiplier 720 is supplied to a first input of an adder 730.

The adder 730 has a second input as well as an output. Connected between its output and its second input is a latch 750. Both the adder 730 and the latch 750 are clocked by the clock signal from the clock signal generator 590. In a customary manner, the adder 730 adds the samples at its two inputs. Since the sum is fed back to its second input via the latch 750, the adder 730 in fact behaves as an accumulator. In addition to being provided to its second input via the latch 750, the output of the adder 730 is provided to a first input ("a") of a divider 740, which also has a second input ("b") as well as an output. The second input ("b") of the divider 740 is the output of the cycle counter 580. The divider 740 is clocked by the clock signal from the clock signal generator 590 and is operable to divide the output of the adder 730 by the output of the counter 580.

Those skilled in the art will therefore appreciate that a differential correlator such as the one just described with reference to Fig. 7 evaluates the following function:

$$5 \quad E[(x(i-s)-y(i))^2],$$

where $y(i)$ is generally either equal to $x(i-s)$ for some value of "s" (when the two signals are truly correlated) or is totally uncorrelated with $x(i)$ (when the two signals are truly uncorrelated). Consequently, \mathcal{R} can be viewed as a normalized version of the "anti-correlation" existing between the input signal sample stream and the switched signal sample stream.

15 It has been found that such "anti-correlation" as computed by the differential correlators 560 leads to particularly robust results in the presence of certain kinds of noise and other system parametric errors or tolerance. As an example, common mode noise disappears from the calculation of the correlation value. As a second example, a 5% amplitude mismatch between the two inputs into the connection verification system 330 results in a 0.25% error component at the output of the multiplier 720.

25 Returning to the description of the connection verification system 330 in Fig. 5, the output of the divider in each of the differential correlators 560A...560E is provided to a minimum detector 570. The minimum detector 570 is adapted to locate the minimum value of " \mathcal{R} " by finding the minimum among the outputs of the differential correlators 560. This minimum anti-correlation value, denoted " \mathcal{R}_{\min} ", corresponds to the minimum anti-correlation value for the current input signal and the current values of p_x , p_y , λ_x and λ_y .

Although the above anti-correlation function has been found to give robust results under certain circumstances, it is within the scope of the invention to use other anti-correlation functions, depending on the operational requirements of the connection verification system 330. Of course, it is also within the scope of the invention to use a correlation function (instead of an anti-correlation function) and to consequently replace the minimum detector 570 with a maximum detector.

Assuming that an anti-correlation system is used, the minimum detector 570 supplies a controller 540 with the minimum anti-correlation value " R_{\min} " for the current input signal (p_x) at the current input wavelength (λ_x) and for the current switched signal (p_y) at the current switched wavelength (λ_y). The controller 540 uses the value of R_{\min} for different values of p_x , p_y , λ_x and λ_y to perform connection verification for the subset of mappings:

$$\{[A_{\lambda}], \lambda \in \{A, \dots, M\}\} \cup \{[C_{\lambda_x}] \times [B_{\lambda_y}], (\lambda_x, \lambda_y) \in \{A, \dots, M\}^2\},$$

which the connection verification system controller 540 receives from the switch controller 150. The connection verification system controller 540 and the switch controller 150 thus communicate, e.g., for directed connectivity verification during switching and to download anticipated cross-connection maps against which the controller 540 can evaluate its measurements.

The controller 540 is adapted to exert control over p_x , p_y , λ_x and λ_y in the following way. Control of the wavelength λ_x of the input signal is achieved by a select signal provided

along select line 542 connected to the first selectors in the input signal front end 320, while control of the switch matrix port px of the input signal is achieved by a select signal provided along select line 544 connected to the second selector in the input signal front end 320. Control of the wavelength λ_y of the switched signal is achieved by a select signal provided along select line 543 connected to the first selectors in the switched signal front end 350, while control of the switch matrix port py of the switched signal is achieved by a select signal provided along select line 546 connected to the second selector in the switched signal front end 350.

In operation, the controller 540 executes a connection verification algorithm, an example of which is now described in greater detail with reference to the flowchart of Fig. 6A, and with continued reference to the subsystem block diagram of Fig. 3.

610: The controller 540 selects $\lambda_x = \lambda_y = \lambda = A$. This value is jointly sent to both front ends 320, 350 along select line 542 and select line 543. The controller 540 proceeds to STEP 612.

612: The controller 540 tests mapping $[A_\lambda]$ for the current value of p (which, in a first instance, will be equal to A). The testing of mapping $[A_\lambda]$ will be described herein below in greater detail with reference to Fig. 6B. The controller 540 proceeds to STEP 614.

614: The controller 540 checks to see whether $\lambda = M$. If $\lambda \neq M$, the controller 540 proceeds to STEP 616;

however, if $\lambda = M$, the controller proceeds to STEP 618.

616: The controller 540 sets λ to the next value in the set $\{A, \dots, M\}$ and returns to STEP 612, where the mapping $[A_\lambda]$ will be tested for this next value of p .

618: All mappings $[A_\lambda]$ have now been tested and it is time for the controller 540 to test each of the mappings $[C_{\lambda x}] \times [B_{\lambda y}]$ for $(\lambda x, \lambda y) \in \{A, \dots, M\}^2$. For notational convenience, the composite mapping $[C_{\lambda x}] \times [B_{\lambda y}]$ will hereinafter be denoted $[G_{\lambda x, \lambda y}]$. At this point, λx is initialized to A and the controller 540 proceeds to STEP 620.

620: Here, λy is also initialized to A , which sets the stage for testing the mapping $[G_{A,A}] = [C_A] \times [B_A]$ in a first instance. The controller then proceeds to STEP 622.

622: The controller tests mapping $[G_{\lambda x, \lambda y}] = [C_{\lambda x}] \times [B_{\lambda y}]$ for the current values of λx and λy (which, in a first instance, are both equal to A). The testing of mapping $[G_{\lambda x, \lambda y}]$ will be described herein below in greater detail with reference to Fig. 6C. After testing $[G_{\lambda x, \lambda y}]$, the controller 540 proceeds to STEP 624.

624: The controller 540 checks to see whether $\lambda y = M$. If $\lambda y \neq M$, the controller 540 proceeds to STEP 626; however, if λy is equal to M , the controller proceeds to STEP 628.

626: The controller 540 sets λ_y to the next value in the set $\{A, \dots, M\}$ and returns to STEP 622, where the mapping $[G_{\lambda_x, \lambda_y}]$ will be tested for this next value of λ_y while retaining the same value of λ_x .

5

628: The controller 540 checks to see whether $\lambda_x = M$. If $\lambda_x \neq M$, the controller 540 proceeds to STEP 630; however, if λ_x is equal to M , the controller 540 has verified all mappings $[G_{\lambda_x, \lambda_y}]$ and therefore returns to the beginning of the connection verification algorithm at STEP 610.

10

630: The controller 540 sets λ_x to the next value in the set $\{A, \dots, M\}$ and returns to STEP 620, where the mapping $[G_{\lambda_x, \lambda_y}]$ will be tested for this next value of λ_x and for all values of λ_y .

15

The above algorithm shows, at a high level, how connection verification is effected for all mappings in the set

20

$$\{[A_\lambda], \lambda \in \{A, \dots, M\}\} \cup$$

$$\{[G_{\lambda_x, \lambda_y}] = [C_{\lambda_x}] \times [B_{\lambda_y}], (\lambda_x, \lambda_y) \in \{A, \dots, M\}^2\}.$$

25 Details regarding the testing of a particular mapping $[A_\lambda]$ are now described with reference to steps 612A-612L in Fig. 6B and details regarding the testing of a particular mapping $[G_{\lambda_x, \lambda_y}]$ are now described with reference to steps 622A-622H in Fig. 6C.

30

612A: The controller 540 selects p_x to be equal to 1. This value is provided to the input signal front end along select line 544, which effectively results in selection of the input signal arriving

at input port 1 of switch matrix 110 λ . The controller 540 proceeds to STEP 612B.

612B: For the chosen value of p_x , the controller 540 selects p_y such that $A_{\lambda}(p_x, p_y) = 1$. In other words, the controller 540 sends a control signal to the switched signal front end 350 along select line 546, where p_y is chosen such that the switched signal passed by the switched signal front end 350 is the one switched signal that is expected to be a delayed version of the input signal currently being passed by the input signal front end 320. The controller 540 proceeds to STEP 612C.

612C: The controller 540 reads the minimum anti-correlation value " R_{\min} " from the minimum detector 570 and proceeds to STEP 612D.

612D: The controller 540 determines whether R_{\min} is close to zero, as would be expected since the switched signal was chosen to be a delayed version of the input signal. If $R_{\min} \approx 0$, the controller 540 proceeds to STEP 612E. However, if R_{\min} is not close to zero, then there is a problem with the connection and the controller 540 proceeds to STEP 612H.

612E: Since $R_{\min} \approx 0$, it is concluded that the connection involving input signal p_x and switched signal y for the current values of p_x and p_y is valid. This condition can be reported to an external module and the controller 540 subsequently proceeds to STEP 612F.

612F: The controller 540 verifies whether $p_x = N$, i.e., whether all input signals have been verified for the current wavelength (the current value of λ). If so, the controller 540 exits STEP 612 altogether and the connection verification algorithm proceeds to STEP 614 (see Fig. 6A). Otherwise, if $p_x \neq N$, then the controller 540 moves to STEP 612G.

612G: The controller 540 increments p_x , following which the controller 540 returns to STEP 612B, where connectivity for the next row in $[A_\lambda]$ is verified.

612H: In order to try and determine where input signal p_x has been mis-connected (or whether the connection involving input signal p_x has been lost altogether), the controller 540 selects, in a deterministic scanning sequence, a different value of y with the constraint that $A_\lambda(p_x, p_y) = 0$. In other words, the controller 540 sends a control signal to the switched signal front 350 end along select line 546, causing the switched signal front end 350 to pass a switched signal p_y which is not expected to be a delayed version of input signal p_x currently being passed by the input signal front end 320. The controller 540 proceeds to STEP 612I.

612I: The controller 540 reads the minimum anti-correlation value " R_{\min} " from the minimum detector 570 and proceeds to STEP 612J.

612J: Although it is expected that R_{\min} will not be close to zero (i.e., $R_{\min} \gg 0$), it may happen that $R_{\min} \approx 0$, for example when a mis-connection has taken place. Accordingly, the controller 540 determines

whether R_{\min} is not close to zero. If R_{\min} is not close to zero, there is no mis-connection for the current values of p_x and p_y and the controller 540 proceeds to STEP 612L. However, if R_{\min} is close to zero, then correlation has been detected for input signal x and switched signal y and the controller 540 proceeds to STEP 612K.

612K: Here, the controller 540 concludes that a mis-connection has taken place and may report this conclusion to an external module. The controller 540 subsequently proceeds to STEP 612F, previously described. However, before proceeding to STEP 612F, the controller 540 may "camp" on this combination of p_x and p_y for a little while longer in order to ascertain, with greater confidence, that a mis-connection has occurred between input signal p_x and switched signal p_y for the current values of p_x and p_y .

612L: The controller 540 checks to see if an attempt has been made to correlate all switched signals (values of p_y) with input signal p_x . If not, then a mis-connection is still possible and the controller returns to STEP 612H. However, if there are no more values of p_y which would lead to $A_{\lambda}(p_x, p_y) = 0$ for the current value of p_x , then the controller proceeds to STEP 612M.

612M: As there are no more possibilities for detecting a mis-connection, the controller 540 provisionally concludes that the connection involving input signal p_x has been "lost". The provisional conclusion of a "lost" connection can be reported

to an external module, after which the controller 540 proceeds to STEP 612F, previously described. It should be noted that the conclusion of a "lost" connection is provisional because a connection may appear to be "lost" although it may in fact be going by $[C_{\lambda x}] \times [B_{\lambda}]$ (for $\lambda x \neq \lambda$) and not straight through $[A_{\lambda}]$ and therefore further testing of composite mappings will have to be performed (see description of STEP 622 below).

With reference now to Fig. 6C, the sub-steps of STEP 622 executed in association with the testing of a particular mapping $[G_{\lambda x, \lambda y}] = [C_{\lambda x}] \times [B_{\lambda y}]$ are now described.

622A: The controller 540 selects p_x to be equal to 1. This value is provided to the input signal front end along select line 544, which effectively results in selection of the input signal arriving at input port 1 of switch matrix 110 λx (where λx and λy were selected prior to entering STEP 622 to begin with - see Fig. 6A). The controller 540 proceeds to STEP 622B.

622B: For the chosen value of p_x , the controller 540 selects p_y such that $G_{\lambda x, \lambda y}(p_x, p_y) = 1$. In other words, the controller 540 sends a control signal to the switched signal front end 350 along select line 546, effectively resulting in selection of the switched signal at output port p_y of optical switch matrix 110 λy . Here, p_y is chosen such that the switched signal passed by the switched signal front end 350 is the one switched signal that is expected to correspond to a delayed version of the input signal currently being passed by the input signal

front end 320. The controller 540 proceeds to STEP 622C.

5 **622C:** The controller 540 reads the minimum anti-correlation value " R_{min} " from the minimum detector 570 and proceeds to STEP 622D.

10 **622D:** The controller 540 determines whether R_{min} is close to zero, as would be expected since the switched signal was chosen to be a delayed version of the input signal. If $R_{min} \approx 0$, the controller 540 proceeds to STEP 622E. However, if R_{min} is not close to zero, then there is a problem with the connection and the controller 540 proceeds to STEP 622H.

15 **622E:** Since $R_{min} \approx 0$, it is concluded that the connection involving input signal p_x and switched signal p_y for the current values of p_x and p_y is valid. This condition can be reported to an external module and the controller 540 subsequently proceeds to STEP 622F.

20 **622F:** The controller 540 verifies whether $p_x = N$, i.e., whether all input signals have been verified for the current values of λ_x and λ_y . If so, the controller 540 exits STEP 622 altogether and the connection verification algorithm proceeds to STEP 624 (see Fig. 6A). Otherwise, if $p_x \neq N$, then the controller 540 moves to STEP 622G.

25 **622G:** The controller 540 increments p_x , following which the controller 540 returns to STEP 622B, where

connectivity for the next row in $[G_{\lambda x, \lambda y}]$ is verified.

622H:

In order to try and determine where input signal p_x has been mis-connected (or whether the connection involving input signal p_x has been lost altogether), the controller 540 selects, in a deterministic scanning sequence, a different value of p_y with the constraint that $G_{\lambda x, \lambda y}(x, y) = 0$. In other words, the controller 540 sends a control signal to the switched signal front 350 end along select line 546, causing the switched signal front end 350 to pass a switched signal p_y which is not expected to be a delayed version of input signal p_x currently being passed by the input signal front end 320. The controller 540 proceeds to STEP 622I.

622I:

The controller 540 reads the minimum anti-correlation value " R_{\min} " from the minimum detector 570 and proceeds to STEP 622J.

622J:

Although it is expected that R_{\min} will not be close to zero (i.e., $R_{\min} \gg 0$), it may happen that $R_{\min} \approx 0$, for example when a mis-connection has taken place. Accordingly, the controller 540 determines whether R_{\min} is not close to zero. If R_{\min} is not close to zero, there is no mis-connection for the current values of p_x and p_y and the controller 540 proceeds to STEP 622L. However, if R_{\min} is close to zero, then correlation has been detected for input signal p_x and switched signal p_y and the controller 540 proceeds to STEP 622K.

622K: Here, the controller 540 concludes that a mis-connection has taken place and may report this conclusion to an external module. This conclusion may override a previously held conclusion of "lost connection", which may have been found at STEP 612M (see Fig. 6B). The controller 540 subsequently proceeds to STEP 622F, previously described. Before proceeding to STEP 622F, the controller 540 may "camp" on this combination of px and py for a little while longer in order to ascertain, with greater confidence, that a mis-connection has occurred between input signal px and switched signal py for the current values of px and py.

622L: The controller 540 checks to see if an attempt has been made to correlate all switched signals (values of py) with input signal px. If not, then a mis-connection is still possible and the controller returns to STEP 622H. However, if there are no more values of py which would lead to $A_{\lambda}(px, py) = 0$ for the current value of px, then the controller proceeds to STEP 622M.

622M: As there are no more possibilities for detecting a mis-connection, the controller 540 definitively concludes that the connection involving input signal px has been "lost". This conclusion can be reported to an external module, after which the controller 540 proceeds to STEP 622F, previously described.

Once the integrity of each connection has been checked, the results can be analyzed to see whether there are any malfunctions or mis-connections, and, if so, whether such

malfunctions or mis-connections are due to a particular optical switch matrix or a particular mirror on a particular optical switch matrix or a particular input or output port on a particular optical switch matrix, etc.

5

Thus, malfunctions and mis-connections which are caused by the straight-through mapping $[A_\lambda]$ of a particular optical switch matrix 110λ will be "caught" using the above algorithm. Also, it will be possible to identify when a mapping $[G_{\lambda x, \lambda y}]$ has failed, although isolation of the error to either $[C_{\lambda x}]$ or $[B_{\lambda y}]$ may require further investigation (see subsystems described later with reference to Figs. 12 and 14).

Those skilled in the art should appreciate that other embodiments of the connection verification system 330 are possible. For example, Fig. 18 shows a connection verification system 330' wherein the delay "resolution" (i.e., the increment in the relative delay between adjacent differential correlators) is equal to the delay produced by one of the delay elements 550, 552. This delay resolution is twice as fine as the delay resolution of the connection verification 330 previously described with reference to Fig. 5. The finer resolution is obtained by supplementing the differential correlators 560A...560E of Fig. 5 with additional differential correlators 560F, 560G and 560H and 560I.

Specifically, differential correlator 560F is fed an undelayed sample from the input signal tapped delay line and a three times delayed sample from the switched signal tapped delay line (for a relative delay of -3), differential correlator 560G is fed a once delayed sample from the input signal tapped delay line and a twice delayed sample from the switched signal tapped delay line (for a relative delay of -

1), differential correlator 560H is fed a twice delayed sample from the input signal tapped delay line and a once sample from the switched signal tapped delay line (for a relative delay of +1) and differential correlator 560I is fed
5 a three times delayed sample from the input signal tapped delay line and an undelayed sample from the switched signal tapped delay line (for a relative delay of +3). Thus, it is seen that all relative delays between -4 and +4 unit delays are "caught" by the bank of differential correlators
10 560A...560I, the outputs of which are fed to an augmented minimum detector 570'.

Those skilled in the art should also appreciate that it is within the scope of the invention to provide a connection
15 verification system in which parallel correlation operations are performed on multiple pairs of input and switched signals. Thus, multiple sets of λ_x and p_x and multiple sets of λ_y and p_y could be generated by the controller 540 at any one time.

20 It should further be appreciated that many variations of the front ends 320, 350 are within the scope of the present invention. For example, suitable variations of the front ends 320, 350 are described in provisional United States
25 Patent Application Serial No. 60/207,292 to Graves et al., entitled "Optical Switch with Connection Verification", filed May 30, 2000 and hereby incorporated by reference herein.

Other suitable variations of the front ends 320, 350 are
30 shown in Figs. 8 through 11. With reference to Fig. 8, there is shown an alternative embodiment of a front end 810, comprising a MEMS-based optical switch or similar optical switch matrix 820 for receiving N WDM signals from the N splitters 310A...310N. In response to the signal on select

line 544, the optical switch matrix 820 selectably passes one of the received WDM signals through to a single output port. The output port of the optical switch matrix 820 is connected to a WDD device 830, which splits the selected WDM signal into its M wavelength components. Each of M output ports of the WDD device 830 carries a single-carrier optical signal which is fed to a respective one of a plurality of receivers 840. Each of the receivers 840 may have a wide optical bandwidth or it may be tuned to the specific wavelength of the single-carrier optical signal received from the WDD device 830. Each receiver operates to convert the incoming single-carrier optical signal from its specific associated WDD port into a band-limited electrical signal. The output of each receiver is fed to a respective input of a selector 850. Thus, the selector 850 receives M band-limited electrical signals, one for each wavelength in the system. In accordance with a select signal received along select line 542, the selector 850 controllably passes one of these M band-limited electrical signals to its output.

Another alternative embodiment of a front end is shown in Fig. 9, where the front end 910 comprises an optical switch matrix 920 for receiving N WDM signals from the N splitters 310A...310N. In response to the signal on select line 544, the optical switch matrix 920 selectably passes one of the N received WDM signals through to a single output port. The output port of the optical switch matrix 920 is connected to a WDD device 930, which splits the selected WDM signal into its M wavelength components. Each of M output ports of the WDD device 930 carries a single-carrier optical signal which is fed to a respective input of another optical switch matrix 940. In response to the signal on select line 542, the optical switch matrix 940 selectably passes one of the M received single-carrier optical signals through to a single

output port, which is connected to a receiver 950. The receiver 950 should have a wide optical bandwidth so that it is capable of accommodating a single-carrier optical signal on any one of the M wavelengths. The receiver 950 operates to convert the incoming single-carrier optical signal into a band-limited electrical signal which is fed to the connection verification system 330.

Fig. 10 shows yet another implementation of a front end 1010, in which the WDM signals received from the plurality of splitters 310A...310N are fed to a respective plurality of tunable optical filters 1020. Each tunable optical filter 1020 can be of standard design and has a pass band that is controllable via select line 542. The output of each of the tunable optical filters 1020 is supplied to a respective one of a plurality of receivers 1030. The receivers 1030 should have a wide optical bandwidth so that they are capable of accommodating single-carrier optical signals on any of the M wavelengths in the system. The receivers 1030 operate to convert the incoming single-carrier optical signals into band-limited electrical signals which are fed to respective inputs of a selector 1040. The selector 1040 controllably passes one of these N band-limited electrical signals to its output, in accordance with a select signal received along select line 544.

Fig. 11 shows still another implementation of a front end 1110. Here, an optical switch matrix 1120 is used to receive N WDM signals from the N splitters 310A...310N. In response to the signal on select line 544, the optical switch matrix 1120 selectably passes one of the N received WDM signals through to a single output port. The output port of the optical switch matrix 1120 is connected to a tunable optical filter 1130, which has a pass band that is controllable via select

line 542. The output of the tunable optical filter 1130 is connected to a receiver 1140, which should have a wide optical bandwidth so that it may accommodate a single-carrier optical signal on any of the M wavelengths in the system.

- 5 The receiver 1140 operates to convert the incoming single-carrier optical signal into a band-limited electrical signal which is provided to the connection verification system 330.

10 As alluded to herein above, the connection verification algorithm described in Figs. 6A-6B may determine that an error has occurred in a mapping $[G_{\lambda x, \lambda y}] = [C_{\lambda x}] \times [B_{\lambda y}]$, although it is not possible, from that algorithm, to determine whether the error is in the execution of mapping $[C_{\lambda x}]$ or mapping $[B_{\lambda y}]$. Thus, further investigation is required. Accordingly, Fig. 12 shows a subsystem for verifying whether the mappings $[B_{\lambda}]$ are properly executed by optical switch matrices 110 λ (for $\lambda \in \{A, \dots, M\}$) and Fig. 14 shows a subsystem for verifying whether the mappings $[C_{\lambda}]$ are being properly executed by optical switch matrices 110 λ (for $\lambda \in \{A, \dots, M\}$).

25 With reference first to Fig. 12, there is provided a first set of optical splitters 1210AA...1210AK, 1210BA...1210BK, ..., 1210MA...1210MK, placed in the optical paths of the M•K single-carrier optical signals travelling between the outputs of the wavelength converting switch 120 and inputs N+1...N+K of each of the optical switch matrices 110A...110M. Each of the splitters 1210 diverts a small fraction of the corresponding single-carrier optical signal towards a common input signal front end 1220.

30 With additional reference to Fig. 13, there is shown a detailed block diagram of a possible implementation of the input signal front end 1220. The input signal front end 1220

comprises a plurality of receivers 1310, each of which is operable to convert the corresponding incoming single-carrier optical signal into a band-limited electrical signal. Each of the receivers 1310 can be tuned to the wavelength of the corresponding single-carrier optical signal. Alternatively, the receivers 1310 can have a wide optical bandwidth. The electrical bandwidth of the receivers 1310 either sets a coarse upper bound on the bandwidth of the input signal that is used for correlation purposes within a connection verification system 1230 if that front end is followed by a precise anti-alias filter or may be a precisely determined bandwidth, in which case the additional anti-alias filter is not required. As with the front end of Fig. 4, an example of a suitable electrical bandwidth for the receivers 1310 is 100 MHz, although other higher and lower bandwidths can be used, depending on the operational requirements of the invention.

The front end 1220 also comprises a set of N first selectors 1320A...1320N, each of which receives the output of M respective receivers 1310. Specifically, the first selectors 1320A...1320N are arranged so that first selector 1320n ($n \in \{A, \dots, N\}$) receives those M electrical signals that correspond to the M optical signals provided by optical splitters 1210nA, 1210nB, ..., 1210nM. Operation of the first selectors 1320A...1320N is jointly controlled via select line 542. The output of each of the N first selectors 1320A...1320N is provided to a respective input of a single second selector 1330. Operation of the second selector 1330 is controlled via select line 544. The output of the second selector 1330 is a band-limited electrical signal provided to the input signal side of the connection verification system 1230.

Continuing with the description of the subsystem in Fig. 12, there is also provided a second set of optical splitters

340A...340N, one in the optical path of each switched WDM signal exiting a respective one of the WDM devices 140A...140N. Each of the N optical splitters 340A...340N diverts a small fraction of the corresponding switched WDM signal towards a common switched signal front end 350. The switched signal front end 350 could be designed as any of the front ends 410, 810, 910, 1010, 1110 in Figs. 4, 8, 9, 10, 11, respectively. The switched signal front end 350 provides a band-limited electrical signal to the switched signal side of the connection verification system 1230.

Thus, it is seen that the connection verification system 1230 receives a band-limited input signal from the input signal front end 1220 and also receives a band-limited switched signal from the switched signal front end 350. In addition, the connection verification system 1230 accepts the set of mappings $[B_\lambda]$, $\lambda \in \{A, \dots, M\}$, from the switch controller 150. The connection verification system 1230 processes its inputs and produces results indicative of whether each input signal that is intended to be switched from the wavelength converting switch 120 directly out of the switching core by one of the optical switch matrices 110A...110M has indeed been properly switched.

To this end, the connection verification system 1230 could be implemented virtually identically to the connection verification system 330 of Fig. 5. The only fundamental difference is that the controller in the connection verification system 1230 would run a slightly different connection verification algorithm, wherein only steps 610-614 would be required and wherein references to $A_\lambda(px, py)$, $1 \leq px, py \leq N$ are replaced with references to $B_\lambda(px, py)$, $1 \leq px \leq K$, $1 \leq py \leq N$.

An alternative embodiment of the subsystem of Fig. 12 is shown in Fig. 12A, wherein the outputs of the optical splitters 1210AA...1210AK, 1210BA...1210BK, ..., 1210MA...1210MK are regrouped by a plurality of WDM devices 1280A...1280K.

5 Specifically, the output of optical splitters 1210AA, 1210BA, ..., 1210MA is provided to WDD device 1280A, the output of optical splitters 1210AB, 1210BB, ..., 1210MB is provided to WDD device 1280B, etc.

10 In this case, only K inputs are provided to the input signal front end, which can hence be made identical to any of the front ends 810, 910, 1010 and 1110 previously described with reference to Figs. 8, 9, 10 and 11, respectively. This alternative embodiment, which requires K additional M-way WDD
15 devices with respect to the embodiment of Fig. 12, affords a significant reduction in the fiber interconnect linking the switching core to the connection verification system.

Fig. 14 shows a subsystem for verifying whether the set of mappings $[C_{\lambda}]$ is being executed properly by the optical
20 switch matrices 110A...110M. Specifically, there is provided a first set of optical splitters 310A...310N, one in the optical path of each input WDM signal exiting a respective one of the WDD devices 130A...130N. Each of the N optical splitters
25 310A...310N diverts a small fraction of the corresponding input WDM signal towards a common input signal front end 320. The input signal front end 320 could be designed as any of the front ends 410, 810, 910, 1010, 1110 in Figs. 4, 8, 9, 10, 11, respectively. The input signal front end 320 provides a
30 band-limited electrical signal to the input signal side of a connection verification system 1430.

In addition, a second set of optical splitters 1440AA...1440AK, 1440BA...1440BK, ..., 1440MA...1440MK are placed in the optical

paths of the $M \times K$ single-carrier optical signals travelling between outputs $N+1 \dots N+K$ of each of the optical switch matrices 110A...110M and the inputs of the wavelength converting switch 120. Each of the splitters 1440 diverts a small fraction of the corresponding single-carrier optical signal towards a common switched signal front end 1450. The switched signal front end 1450 could be implemented as the input signal front end 1220 of Fig. 12, except that select line 544 (which supplies p_x) would be replaced by select line 546 (which supplies p_y). The output of the switched signal front end 1450 is a band-limited electrical signal provided to the switched signal side of the connection verification system 1430.

Aside from accepting a band-limited input signal from the input signal front end 320 and a band-limited switched signal from the switched signal front end 1450, the connection verification system 1430 also accepts the set of mappings $[C_\lambda]$, $\lambda \in \{A, \dots, M\}$, from the switch controller 150. The connection verification system 1430 processes its inputs and produces results indicative of whether each input signal that is intended to be switched out of the switching core and into the wavelength converting switch 120 by one of the optical switch matrices 110A...110M has indeed been properly switched. To this end, the connection verification system 1430 could be implemented virtually identically to the connection verification system 330 of Fig. 5, except that the controller in the connection verification system 1430 would run a slightly different connection verification algorithm, wherein references to $A_\lambda(p_x, p_y)$, $1 \leq p_x, p_y \leq N$ are replaced with references to $C_\lambda(p_x, p_y)$, $1 \leq p_x \leq N$, $1 \leq p_y \leq K$.

An alternative embodiment of the subsystem of Fig. 14 is shown in Fig. 14A, wherein the outputs of the optical

splitters 1440AA...1440AK, 1440BA...1440BK, ..., 1440MA...1440MK are regrouped by a plurality of WDM devices 1490A...1490K. Specifically, the output of optical splitters 1440AA, 1440BA, ..., 1440MA is provided to WDD device 1490A, the output of
 5 optical splitters 1440AB, 1440BB, ..., 1440MB is provided to WDD device 1490B, etc.

In this case, only K inputs are provided to the switched signal front end, which can hence be made identical to any of
 10 the front ends 810, 910, 1010 and 1110 previously described with reference to Figs. 8, 9, 10 and 11, respectively. This alternative embodiment, which requires K additional M-way WDD devices with respect to the embodiment of Fig. 14, affords a significant reduction in the fiber interconnect linking the
 15 switching core to the connection verification system.

With reference to Figs. 15, 16 and 17, there are shown subsystems for respectively verifying mappings [D], [E] and [F] of the wavelength converting switch 120. Specifically,
 20 in Fig. 15, optical splitters 1510AA...1510AK, 1510BA...1510BK, ..., 1510MA...1510MK are placed in the optical paths of the M•K single-carrier optical signals travelling between outputs N+1...N+K of each of the optical switch matrices 110A...110M and the inputs of the wavelength converting switch 120. Each of
 25 the splitters 1510 diverts a small fraction of the corresponding single-carrier optical signal towards a common first front end 1520. The first front end 1520 could be implemented as the input signal front end 1220 of Fig. 12. The output of the first front end 1520 is a band-limited
 30 electrical signal provided to a first signal side of a connection verification system 1530.

Also provided in the subsystem of Fig. 15 are optical splitters 1540AA...1540AK, 1540BA...1540BK, ..., 1540MA...1540MK,

which are placed in the optical paths of the $M \cdot K$ single-carrier optical signals travelling between the outputs of the wavelength converting switch 120 and inputs $N+1 \dots N+K$ of each of the optical switch matrices 110A...110M. Each of the
 5 splitters 1540 diverts a small fraction of the corresponding single-carrier optical signal towards a common second front end 1550. The second front end 1550 could be implemented as the input signal front end 1220 of Fig. 12, except that select line 544 (which supplies p_x) would be replaced by
 10 select line 546 (which supplies p_y). The output of the second front end 1550 is a band-limited electrical signal provided to a second signal side of the connection verification system 1530.

15 The connection verification system 1530 also accepts mapping [D] from the switch controller 150. The connection verification system 1530 processes its inputs and produces results indicative of whether each input signal that is intended to be switched by the wavelength converting switch
 20 120 in accordance with mapping [D] has indeed been properly switched. To this end, the connection verification system 1530 could be implemented virtually identically to the connection verification system 330 of Fig. 5, except that the controller in the connection verification system 1530 would
 25 run a slightly different connection verification algorithm, wherein references to $A_\lambda(p_x, p_y)$, $1 \leq p_x, p_y \leq N$ are replaced with references to $D(((\lambda-1) \cdot K) + p_x, ((\lambda-1) \cdot K) + p_y)$, $1 \leq p_x, p_y \leq K$, since [D] is an $M \cdot K \times M \cdot K$ matrix and each $[A_\lambda]$ is an $N \times N$ matrix.

30 An alternative embodiment of the subsystem of Fig. 15 is shown in Fig. 15A, wherein the outputs of the optical splitters 1510AA...1510AK, 1510BA...1510BK, ..., 1510MA...1510MK are regrouped by a plurality of WDM devices 1280A...1280K.

Specifically, the output of optical splitters 1510AA, 1510BA, ..., 1510MA is provided to WDD device 1280A, the output of optical splitters 1510AB, 1510BB, ..., 1510MB is provided to WDD device 1280B, etc. In addition, there is provided a plurality of WDM devices 1490A...1490K for regrouping the outputs of optical splitters 1540AA, 1540BA, ..., 1540MA in a similar fashion.

In this case, only K inputs are provided to each of the front ends, which can hence be made identical to any of the front ends 810, 910, 1010 and 1110 previously described with reference to Figs. 8, 9, 10 and 11, respectively. This alternative embodiment, which requires $2 \cdot K$ additional M-way WDD devices with respect to the embodiment of Fig. 15, affords a significant reduction in the fiber interconnect linking the switching core to the connection verification system.

Fig. 16 is similar to Fig. 15 in that it retains the optical splitters 1540AA...1540AK, 1540BA...1540BK, ..., 1540MA...1540MK, and the second front end 1550. However, the subsystem of Fig. 16 includes a plurality of optical splitters 1610A...1610R, which are placed in the optical paths of the R add carriers 180A...180R. Each of the splitters 1610 diverts a small fraction of the corresponding single-carrier optical signal towards a common first front end 1620, which could be implemented as a scaled version of the second front end 1550. The band-limited electrical signals output by the first and second front ends 1620, 1550 are fed to a connection verification system 1630, which also accepts mapping [E] from the switch controller 150.

The connection verification system 1630 processes its inputs and produces results indicative of whether each add carrier

that is intended to be switched by the wavelength converting switch 120 in accordance with mapping [E] has indeed been properly switched. To this end, the connection verification system 1630 could be implemented virtually identically to the connection verification system 330 of Fig. 5, except that the controller in the connection verification system 1630 would run a slightly different connection verification algorithm, wherein references to $A_{\lambda}(p_x, p_y)$, $1 \leq p_x, p_y \leq N$ are replaced with references to $E(p_x, ((\lambda-1) \cdot K) + p_y)$, $1 \leq p_x \leq R$, $1 \leq p_y \leq K$, since [E] is an $R \times M \cdot K$ matrix and each $[A_{\lambda}]$ is an $N \times N$ matrix.

An alternative embodiment of the subsystem of Fig. 16 is shown in Fig. 16A, wherein the outputs of the optical splitters 1540AA...1540AK, 1540BA...1540BK, ..., 1540MA...1540MK are regrouped by a plurality of WDM devices 1490A...1490K. Specifically, the output of optical splitters 1540AA, 1540BA, ..., 1540MA is provided to WDD device 1490A, the output of optical splitters 1540AB, 1540BB, ..., 1540MB is provided to WDD device 1490B, etc.

In this case, only K inputs are provided to the front end (formerly 1550 in Fig. 16), which can hence be made identical to any of the front ends 810, 910, 1010 and 1110 previously described with reference to Figs. 8, 9, 10 and 11, respectively. This alternative embodiment, which requires K additional M-way WDD devices with respect to the embodiment of Fig. 16, affords a significant reduction in the fiber interconnect linking the switching core to the connection verification system.

With reference now to Fig. 17, there is shown a subsystem which is similar to that of Fig. 15 in that it retains the optical splitters 1510AA...1510AK, 1510BA...1510BK, ...,

1510MA...1510MK, and the first front end 1520. However, the subsystem of Fig. 17 includes a plurality of optical splitters 1710A...1710R, which are placed in the optical paths of the R drop carriers 170A...170R. Each of the splitters 1710
 5 diverts a small fraction of the corresponding single-carrier optical signal towards a common second front end 1750, which could be implemented as a scaled version of the first front end 1520. The band-limited electrical signals output by the first and second front ends 1520, 1750 are fed to a
 10 connection verification system 1730, which also accepts mapping [F] from the switch controller 150.

The connection verification system 1730 processes its inputs and produces results indicative of whether each of the
 15 signals arriving from the optical switch matrices 110A-110M and intended to be switched towards one of the drop carriers 170A...170R in accordance with mapping [F] has indeed been properly switched. To this end, the connection verification system 1730 could be implemented virtually identically to the
 20 connection verification system 330 of Fig. 5, except that the controller in the connection verification system 1730 would run a slightly different connection verification algorithm, wherein references to $A_{\lambda}(x,y)$, $1 \leq p_x, p_y \leq N$ are replaced with references to $E(((\lambda-1) \cdot K) + p_x, p_y)$, $1 \leq p_x \leq K$, $1 \leq p_y \leq$
 25 R , since [E] is an $M \cdot K \times R$ matrix and each $[A_{\lambda}]$ is an $N \times N$ matrix.

An alternative embodiment of the subsystem of Fig. 17 is shown in Fig. 17A, wherein the outputs of the optical
 30 splitters 1510AA...1510AK, 1510BA...1510BK, ..., 1510MA...1510MK are regrouped by a plurality of WDM devices 1280A...1280K. Specifically, the output of optical splitters 1510AA, 1510BA, ..., 1510MA is provided to WDD device 1280A, the output of

optical splitters 1510AB, 1510BB, ..., 1510MB is provided to WDD device 1280B, etc.

5 In this case, only K inputs are provided to the front end (formerly 1520 in Fig. 17), which can hence be made identical to any of the front ends 810, 910, 1010 and 1110 previously described with reference to Figs. 8, 9, 10 and 11, respectively. This alternative embodiment, which requires K additional M-way WDD devices with respect to the embodiment
10 of Fig. 17, affords a significant reduction in the fiber interconnect linking the switching core to the connection verification system.

15 It should be understood that all or part of the input signal front end and/or the switched signal front end in any of the above described embodiments may be shared with a power spectrum flattening system such as that described in United States Patent Application Serial No. 09/580,495 to Graves et al., filed May 30, 2000, entitled "Optical Switch with Power
20 Equalization", assigned to the assignee of the present invention and hereby incorporated by reference herein.

It should further be appreciated that in some embodiments of the invention, all or part of the functionality previously
25 described herein with respect to the controller 540 and the differential correlators 560A...560G may be implemented as pre-programmed hardware or firmware elements (e.g., application specific integrated circuits (ASICs), electrically erasable programmable read-only memories (EEPROMs), etc.), or other
30 related components.

In other embodiments of the invention, all or part of the functionality previously described herein with respect to the controller 540 and the differential correlators 560A...560G may

be implemented as software consisting of a series of instructions for execution by a computer system. The series of instructions may be written in a number of programming languages for use with many computer architectures or operating systems. For example, some embodiments may be implemented in a procedural programming language (e.g., "C") or an object oriented programming language (e.g., "C++" or "JAVA").

10 The series of instructions could be stored on a medium which is fixed, tangible and readable directly by the computer system, (e.g., removable diskette, CD-ROM, ROM, or fixed disk), or the instructions could be stored remotely but transmittable to the computer system via a modem or other interface device (e.g., a communications adapter) connected to a network over a transmission medium. The transmission medium may be either a tangible medium (e.g., optical or analog communications lines) or a medium implemented using wireless techniques (e.g., microwave, infrared or other transmission schemes).

While specific embodiments of the present invention have been described and illustrated, it will be apparent to those skilled in the art that numerous modifications and variations can be made without departing from the scope of the invention as defined in the appended claims.